arXiv:1511.06436v1 [cs.SC] 19 Nov 2015

On the robust hardness of Gröbner basis computation

Gwen Spencer

Smith College, Northampton, MA

David Rolnick

Massachusetts Institute of Technology, Cambridge, MA

Abstract

We introduce a new problem in the approximate computation of Gröbner bases that allows the algorithm to ignore a constant fraction of the generators - of the algorithm's choice - then compute a Gröbner basis for the remaining polynomial system. The set ignored is subject to one quite-natural structural constraint. For lexicographic orders, when the discarded fraction is less than $(1/4 - \epsilon)$, for $\epsilon > 0$, we prove that this problem cannot be solved in polynomial time, even when the original polynomial system has maximum degree 3 and each polynomial contains at most 3 variables. Qualitatively, even for sparse systems composed of low-degree polynomials, we show that Gröbner basis computation is robustly hard: even producing a Gröbner basis for a large subset of the generators is NP-hard.

Key words: Gröbner Basis, Complexity, Approximation, Hardness, Combinatorial Optimization, Satisfiability

1. Introduction

A classic problem in computational algebra is producing a Gröbner basis for an ideal. Suppose that $\mathbb{K}$ denotes a field. Given a system of polynomials $\mathcal{F} \subseteq \mathbb{K}[x_1, \ldots, x_n]$, let $\langle \mathcal{F} \rangle$ denote the ideal generated by the polynomials of $\mathcal{F}$. For a given term order, a *Gröbner basis* for the ideal $\langle \mathcal{F} \rangle$ is an alternative set of generators for which the leading terms generate the ideal of leading terms in $\langle \mathcal{F} \rangle$.

⋆ D.R. was supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374. Both authors were supported in early stages of the project by the Mathematics Research Communities Program of the AMS.

Email addresses: gwenspencer@gmail.com (Gwen Spencer), drolnick@mit.edu (David Rolnick).

Given a Gröbner basis for an ideal, important problems like the ideal-membership test (and others) can easily be resolved. For more background, see the widely-read textbook of Cox, Little, and O'Shea [6]. The best-known general algorithm for computing a Gröbner basis is due to Buchberger [3] and does not run in polynomial time. Further, the existence of a polynomial-time algorithm for computing Gröbner bases is precluded by a number of hardness results. For general $\mathcal{F}$, the problem is EXPSPACE-complete [18, 20]. The maximum degree of polynomials in a Gröbner basis can also grow exponentially, even for zero-dimensional ideals [9, 21, 22, 23].

Despite strong hardness results for the general problem of computing a Gröbner basis, there is hope for efficient methods for restricted classes of polynomial systems. De Loera et al. [7] give a polynomial-time algorithm for Gröbner basis computation in the case where $\mathcal{F}$ encodes the problem of properly coloring a chordal graph. Indeed, recent work of Cifuentes and Parrilo [5] in computational algebraic geometry has begun to unearth rich connections between restrictions on the graphical representation of the structure of a system of polynomials and efficient algorithmic methods for Gröbner basis computation.

Where is the boundary between polynomial systems for which the highly-useful tool of a Gröbner basis can be efficiently produced, and those systems for which this is impossible? To better understand this boundary, we examine the robust hardness of Gröbner basis computation, leveraging results from combinatorial optimization.

Consider two questions that seek to understand the hardness of a problem:

- How much time and space are required to solve the problem accurately?
- How well can the problem be solved in limited time and space? *In particular, what is the best "approximate solution" that can be guaranteed in polynomial time?*

In combinatorial optimization, both questions are vigorously studied for many famous problems (Graph Coloring, Traveling Salesman, Satisfiability, Maximum Cut, Steiner Tree, etc). The second question is of interest for problems that are "robustly hard": problems which are not only *hard to solve exactly*, but which are also *hard to solve even approximately*. Regarding Gröbner basis computation, much study has been devoted to the first question, but apparently little attention has been devoted to the second. We suspect the difference stems from how naturally-parametrizable the notion of an approximate solution is for optimization problems. [1]

Given a polynomial system $\mathcal{F}$ for which we would like to compute a Gröbner basis, what could it mean to give an approximate answer? De Loera et al. [7] define the *Strong c-Partial Gröbner Problem* which allows the algorithm to ignore c independent sets of variables (and all generators containing them), where an *independent set* of variables is a set such that no pair co-occur in any generator. The algorithm need only return a Gröbner basis for the ideal corresponding to the remaining generators. De Loera et al. show that the *Strong c-Partial Gröbner Problem* problem is NP-hard whenever $c \leq 2\lfloor \frac{k}{3} \rfloor - 1$ (here k is a degree bound for $\mathcal{F}$, and the result holds for any $k \geq 3$). Their proof leverages a result of Khanna, et al. [17] on the NP-hardness of producing a suboptimal proper coloring even when the chromatic number is given.

[1] For example, an *β-approximation algorithm* is defined as a polynomial-time algorithm guaranteed to return a feasible solution whose objective is within a multiplicative factor β of the objective for the best feasible solution.

In this paper, we show that the freedom to selectively ignore nearly a full quarter of the generators does not allow us to compute a Gröbner basis in polynomial time. Formally, we introduce the *q-Fractional Gröbner Basis Problem*, in which the algorithm can ignore a *constant fraction* q of the generators (subject to one natural structural constraint), and return a Gröbner basis for the ideal corresponding to the remaining generators. In Section 2 we prove this problem cannot be solved in polynomial time for $q = 3/4 + \epsilon$, even in severely restricted cases (when each polynomial in $\mathcal{F}$ contains at most 3 variables and has maximum degree 3). For a version with no structural constraints on the ignored set of generators, a weaker-parameter version follows from the same argument (again this holds even in severely restricted cases).

In Section 3, we consider a similar notion of "approximate solution" for graph-coloring, where we are allowed to ignore a certain fraction of the input. We show that either positive or negative hardness results for this problem would have powerful consequences for more traditional models of approximate graph coloring. Finally, in Section 4 we give a new result for the Strong c-Partial Gröbner Problem posed in De Loera, et al. [7]. Rather than focusing on results for low-degree-bound systems, we explore the NP-hardness of the Strong c-Partial Gröbner Problem when c is a large linear function of the degree bound for $\mathcal{F}$.

Our complexity results for the *q-fractional* problems rely on both inapproximability and positive algorithmic results for MAX-3SAT and for graph-coloring. In the last section, we consider the problem of ignoring independent sets of variables as in [7]: our result relies on the NP-hardness of a certain chromatic-number decision problem as shown by Lund and Yannakakis [19].

2. A new flavor of robust hardness: The q-Fractional Gröbner Basis Problem

Given a set of polynomials $\mathcal{F}$ over a set of variables X: for $X_{out} \subseteq X$, let $X_{in} = X \backslash X_{out}$, let $\mathcal{F}_{out}$ denote the subset of $\mathcal{F}$ whose polynomials contain some variable from X_{out}, and let $\mathcal{F}_{in} = \mathcal{F} \backslash \mathcal{F}_{out}$.

Definition 1 (**q-Fractional Gröbner Basis Problem**)**.** Given as input (i) a set of polynomials $\mathcal{F}$ over variables X and (ii) a value $q \in [0, 1]$, output the following:
- $X_{out} \subseteq X$, such that $|\mathcal{F}_{out}| \leq (1 - q)|\mathcal{F}|$.
- A Gröbner basis for $\mathcal{F}_{in}$ with respect to any lexicographic order.

Notice that for $q = 1$, this is the traditional Gröbner basis problem for $\mathcal{F}$. The set X_{out} corresponds to a set of variables *chosen by the algorithm to be ignored*: all the polynomials containing variables from X_{out} can be ignored, and the algorithm need only compute a Gröbner basis for the remaining set of polynomials $\mathcal{F}_{in}$.

Theorem 2 (**Main Result**)**.** *Assume that we are working over a polynomial ring $\mathbb{K}[x_1, x_2, x_3, \ldots, x_n]$. For any $\epsilon > 0$: there is no polynomial-time algorithm $\mathcal{A}$ that solves the $(3/4 + \epsilon)$-Fractional Gröbner Basis Problem (unless $P = NP$). This statement holds even when $\mathcal{F}$ has maximum degree 3 and each polynomial from $\mathcal{F}$ contains at most 3 variables.*

Our proof of Theorem 2 is by reduction from the well-known MAX-3SAT Problem. Our reduction combines Håstad's celebrated hardness-of-approximation result for satisfiable MAX-3SAT instances [12] with analysis adapted from Johnson's classic randomized (7/8)-approximation algorithm for MAX-3SAT (which may be derandomized with the method of conditional expectations) [13].

Definition 3 (**Maximum 3-Satisfiability (MAX-3SAT)**)**.** We are given as input a set of logical clauses $\mathcal{C}$ over a set of variables $\{y_1, \ldots, y_n\}$. Each clause in C is composed of 3 variables or their negations, linked by disjunction (e.g. a possible clause is $y_3 \vee \neg y_6 \vee y_{11}$). We must output a true/false assignment for $\{y_1, \ldots, y_n\}$ that satisfies the maximum fraction of the clauses in $\mathcal{C}$.

Definition 4. An algorithm $\mathcal{A}$ is called a "β-approximation algorithm" for an optimization problem if, in time polynomial in the size of the input, $\mathcal{A}$ returns a solution whose value is within a β-multiplicative factor of the optimal value for the input.

Further, an instance of MAX-3SAT is called *satisfiable* if there exists an assignment of true/false values for the variables such that every clause in $\mathcal{C}$ is true. So, for a satisfiable instance of MAX-3SAT the optimal value is 1 (all clauses are satisfied).

Theorem 5 (Theorem 6.5 in Håstad [12])**.** *Assuming that P≠NP, and $\delta > 0$, there is no $(7/8 + \delta)$-approximation algorithm for MAX-3SAT (even when instances are guaranteed to be satisfiable).*

To prove our main result, Theorem 2, we assume, for the sake of contradiction, that such an algorithm $\mathcal{A}$ does exist, and show how to construct a polynomial-time algorithm for satisfiable instances of MAX-3SAT that is strictly better than what is allowed by Håstad's result. Crucially, we also use that under a lexicographic order, possession of a Gröbner basis will allow efficient computation of a point in the variety of the corresponding polynomial system by iterative elimination of one variable at a time.

Proof of Theorem 2. For the sake of contradiction, assume that the $\mathcal{A}$ asserted in Theorem 2 does exist with $q = (3/4 + \epsilon)$ for some fixed $\epsilon > 0$. Given an arbitrary satisfiable instance of the MAX-3SAT Problem with clause set $\mathcal{C}$ and variable set $\{y_1, \ldots, y_n\}$ we find a truth assignment of quality forbidden by Theorem 5 in polynomial time as follows. [2]

Define a polynomial system as follows. For each y_i, create a variable x_i. For each clause, create a polynomial as a product of three terms. If y_j appears in positive form in the clause, then include the term $(x_j - 1)$ in the product. If y_j is negated, include the term $(x_j + 1)$ in the product. For example,

$$y_{11} \vee \neg y_6 \vee y_3 \text{ gives the polynomial } (x_{11} - 1)(x_6 + 1)(x_3 - 1).$$

In this way, we obtain a set of $|\mathcal{C}|$ polynomials of maximum total degree at most 3 in $\mathbb{K}[x_1, x_2, x_3, \ldots, x_n]$. Call this set $\mathcal{F}$. Observe that any satisfying assignment for MAX-3SAT can be interpreted as a point in the variety $\mathcal{V}(\langle \mathcal{F} \rangle)$: set $x_i = 1$ if y_i is true, otherwise

[2] Without loss of generality we may assume that no clause contains a variable in both positive and negative form: such clauses are trivially true and may be ignored.

set $x_i = -1$. As our initial 3-SAT Problem was satisfiable, any subset of $\mathcal{F}$ generates an ideal with non-empty variety.

We now run $\mathcal{A}$ on $\mathcal{F}$. Let X_{out} denote the set of variables which $\mathcal{A}$ selects with

$$|\mathcal{F}_{out}| \leq (1 - (3/4 + \epsilon))|\mathcal{F}| = (1/4 - \epsilon)|\mathcal{C}|.$$

Let $\mathcal{C}_{out}$ denote the set of clauses that correspond to the polynomials of $\mathcal{F}_{out}$. Thus, by definition, $|\mathcal{C}_{out}| \leq (1/4 - \epsilon)|\mathcal{C}|$. Denote $\mathcal{C} \backslash \mathcal{C}_{out}$ by $\mathcal{C}_{in}$. Likewise, let Y_{out} be the set of y_i such that $x_i \in X_{out}$.

Thus, $\mathcal{A}$ computes a Gröbner basis $\mathcal{G}$ for the ideal generated by the $(3/4 + \epsilon)|\mathcal{C}|$ polynomials in $\mathcal{C}_{in}$. We claim that from $\mathcal{G}$ we can reconstruct a solution in the variety of $\mathcal{F}_{in}$. Let $X_{in} = \{x_{i_1}, x_{i_2}, \ldots, x_{i_m}\}$. By the Elimination Theorem (Theorem 2 in §3.1 of [6]), the set $\mathcal{G}_k = \mathcal{G} \cap \mathbb{K}[x_{i_k}, x_{i_{k+1}}, \ldots, x_{i_m}]$ is a Gröbner basis for $\mathcal{V}(\mathcal{F}_{in}) \cap \mathbb{K}[x_{i_k}, x_{i_{k+1}}, \ldots, x_{i_m}]$.

Given $a \in \mathcal{V}(\mathcal{F}_{in})$, let $\pi_k(a)$ be the element of $\mathbb{K}[x_{i_k}, x_{i_{k+1}}, \ldots, x_{i_m}]$ formed by projecting a. By the Closure Theorem (Theorem 3 in §3.2 of [6]), the variety of solutions to $\mathcal{G}_k$ is the smallest affine variety containing $\pi_k(a)$ for every $a \in \mathcal{V}(\mathcal{F}_{in})$. Note that by our construction of $\mathcal{F}$, the variety $\mathcal{V}(\mathcal{F}_{in})$ is the union of a finite number of sets V_j, where each V_j is the product of sets of the form $\mathbb{R}$, $\{1\}$, or $\{-1\}$. This is because, if we fix all the variables but one, the remaining variable is either constrained to 1 or -1 or is completely unconstrained.

Hence, every projection of $\mathcal{V}(\mathcal{F}_{in})$ is itself an affine variety; thus, every solution to $\mathcal{G}_k$ extends to a solution to $\mathcal{G}$. Therefore, in order to reconstruct a solution $a \in \mathcal{V}(\mathcal{F}_{in})$, we can iteratively extend a partial solution in $\mathcal{V}(\mathcal{G}_k)$. Specifically, we pick some value of x_{i_k} that solves $\mathcal{G}_k$, given that we have already picked values for $x_{i_{k+1}}, \ldots, x_{i_m}$.

This solution a is a vector of length $|X_{in}|$ which is a mutual zero of all polynomials in $\mathcal{F}_{in}$. Each entry in the vector corresponds to some variable in our original MAX-3SAT instance: if $a_i = 1$, we assign the corresponding variable y_i to be true. If $a_i = -1$, we assign y_i to be false. For values of a_i other than -1 or 1 no assignment is made for y_i. By construction, the fact that a causes every polynomial in $\mathcal{F}_{in}$ to evaluate to 0 means that every clause in $\mathcal{C}_{in}$ evaluates as "true" under the truth assignment constructed so far. Thus far, our constructed truth assignment has assigned truth values only for y_i which correspond to the variables x_i in X_{in} (and gives a perfect truth assignment for the clauses in $\mathcal{C}_{in}$).

The variables y_i in Y_{out} have not yet been assigned truth values. Consider a random procedure that "flips a fair coin" to decide the truth value for each such y_i. By definition, every clause in $\mathcal{C}_{out}$ contains at least one $y_i \in Y_{out}$. For an arbitrary clause $C \in \mathcal{C}_{out}$, the probability that C is true as a result of this random assignment procedure is at least $1/2$ (and may be higher). Thus, letting $\mathcal{C}^*_{out}$ denote the subset of clauses from $\mathcal{C}_{out}$ which are true at the end of the random procedure, and computing the expectation with respect to the random coin flips:

$$E[|\mathcal{C}^*_{out}|] \geq \frac{1}{2}|\mathcal{C}_{out}|.$$

This random assignment method can be derandomized via the method of conditional expectations to get a deterministic assignment for Y_{out} that satisfies at least $|\mathcal{C}_{out}|/2$ clauses. For completeness, a sketch of this standard reasoning follows. [3]

[3] Crucially for our proof, the following argument does not depend on the true maximum fraction of the clauses from $\mathcal{C}_{out}$ that may still be satisfied as a result of our partial truth assignment for $y_i \in Y_{out}$.

For an arbitrarily ordered list of the variables y_i to be randomly assigned, proceed through the list deterministically fixing truth values of one variable at a time as follows: for the top variable y that is not yet assigned, compute and compare the conditional expected values of $E[|\mathcal{C}^*_{out}|]$ given an assignment of true for y and given an assignment of false for y (the expectation is computed given that all variables prior to y have had their truth values fixed deterministically, and that all variables after y will have their truth values decided by fair coin flip). At least one of these conditional expectations is as large as the expected value when y is also decided by fair coin flip (while all variables prior to y have been fixed deterministically), so assign y the truth value with the larger conditional expectation. This procedure gives a deterministic solution of quality at least as high as the original expected value.

Our constructed truth assignment satisfies the $|\mathcal{C}| - |\mathcal{C}_{out}|$ clauses in $\mathcal{C}_{in}$ and at least $|\mathcal{C}_{out}|/2$ of the clauses from $\mathcal{C}_{out}$. By our construction, $|\mathcal{C}_{out}| \leq (1/4 - \epsilon)|\mathcal{C}|$ (where $\epsilon > 0$). Thus, our constructed truth assignment is guaranteed to satisfy more clauses than Theorem 5 permits:

$$\begin{aligned} |\mathcal{C}| - |\mathcal{C}_{out}| + \frac{|\mathcal{C}_{out}|}{2} &= |\mathcal{C}| - \frac{|\mathcal{C}_{out}|}{2} \\ &\geq |\mathcal{C}| - \frac{(1/4 - \epsilon)|\mathcal{C}|}{2} \\ &\geq |\mathcal{C}| \left(1 - \frac{(1/4 - \epsilon)}{2}\right) \\ &\geq |\mathcal{C}| \left(\frac{7}{8} + \frac{\epsilon}{2}\right). \end{aligned}$$

Finally, it is obvious that the reduction described (construction of the $(3/4 + \epsilon)$-*fractional Gröbner basis* instance, and our subsequent deterministic construction of a truth assignment of forbidden quality) is in polynomial time given our initial assumption that $\mathcal{A}$ is a polynomial time algorithm. Thus, unless P=NP, we have arrived at a contradiction. □

Comments. De Loera et al. [7] first noted that the hardness-of-approximation of combinatorial optimization problems would translate to a family of results on the "robust hardness" of Gröbner basis computation in connection with the classical Graph Coloring Problem. Notably, the reduction from graph coloring given in [7] requires more subtle structural conditions on the set of polynomials which may be ignored by an algorithm.

Our Theorem 2 is satisfying because, with only a very natural constraint on the structure of the removed set, a remarkably large set of generators may be removed without resulting in a polynomial-time solvable problem. Furthermore, this robust hardness holds even for very low-degree polynomial systems which are sparse in the sense that each polynomial contains at most 3 variables.

What further restrictions on a polynomial system $\mathcal{F}$ might yield Gröbner Basis Problems or Approximate Gröbner Basis Problems that are solvable in polynomial time? In

If we could guarantee that a large portion of the clauses from $\mathcal{C}_{out}$ could still be satisfied despite the existing partial-truth assignment, our second phase of truth-assignment construction could leverage tighter approximation algorithms for MAX-3SAT.

our proof of Theorem 2 the constructed polynomial system belongs to a highly restricted class: the maximum degree is 3, no variable has degree greater than 1 in any term, a polynomial never contains more than 3 variables, and all coefficients come from the set $\{-1, 0, 1\}$. Qualitatively, a polynomial system $\mathcal{F}$ may be very-highly restricted and yet still yield a $(3/4+\epsilon)$-Fractional Gröbner Basis Problem which cannot be solved efficiently. Next we give a strengthened statement of Theorem 2.

The strengthened version relies on an identical reduction, but we gain a significant additional restriction on $\mathcal{F}$ (without losing anything in our robust hardness result) by using a more recent inapproximability result of Guraswami and Khot [11] for a specialized variant of MAX-3SAT known as "Max Non-Mixed Exactly 3SAT." In this variant, each clause must contain exactly 3 literals, and each clause either has all three literals in positive form or all three literals in negated form (there is "no mixing" of positive-form and negated-form literals within clauses). Guraswami and Khot match Håstad's result for the general case: even this specialized variant is NP-hard to approximate within multiplicative factor better than 7/8, and this is true even for instances guaranteed to be satisfiable. Our reduction works in the same way as before, but the polynomial system constructed from the arbitrary Max Non-Mixed Exactly 3SAT instance has an even simpler form. Thus, even a efficient algorithm which can only solve the $(3/4 + \epsilon)$-Fractional Gröbner Basis Problem for this more restricted class of polynomial systems would allow construction of a truth assignment of forbidden quality.

Theorem 6 (Extension of Theorem 2)**.** *Assume that we are working over a polynomial ring $\mathbb{K}[x_1, x_2, x_3, \ldots, x_n]$. Let $\mathcal{Q}$ denote any class of polynomials that contains all polynomials of the form either*

$$(x_i + 1)(x_j + 1)(x_k + 1) \;\; or \;\; (x_i - 1)(x_j - 1)(x_k - 1)$$

for $i, j, k \in \{1, 2, \ldots, n\}$.

For any $\epsilon > 0$: there is no polynomial-time algorithm $\mathcal{A}$ that solves the $(3/4 + \epsilon)$-Fractional Gröbner Basis Problem with respect to any lexicographic order (unless P=NP). This statement holds even when $\mathcal{F}$ is restricted to contain only polynomials from $\mathcal{Q}$.

Theorem 6 identifies a surprisingly-simple core subset of only $2\binom{n}{3}$ polynomials that are each symmetric in the variables they contain.[4] Restriction of a polynomial system to any class still containing this core is not enough to ensure polynomial-time solvability of even the $(3/4+\epsilon)$-Fractional Gröbner Basis Problem. At a high level, even for severely restricted inputs, we find that robust hardness still holds.

Finally, we close this section by noting that a simplified version of our proof of Theorem 2 also yields a statement that holds with no structural constraint on the set of ignored polynomials. Our q-fractional model allows the algorithm to select a set of variables X_{out}, and ignore all polynomials which contain a variable in X_{out} (as long as this set of polynomials, F_{out}, is not too large). In our reduction this means that every clause corresponding to an ignored polynomial contains some variable that is undecided when coin-flipping begins. This property is essential to the expected quality achieved by the coin-flip-based portion of the truth assignment.

With absolutely no structural constraint on the subset of original generators ignored, a $(1/8 - \epsilon)$-fraction of the generators may be ignored without compromising NP-Hardness:

[4] For geometric intuition on the severity of this restriction consider that a variety defined by such a polynomial is invariant under permutations of the variables appearing in the polynomial.

Theorem 7. *Assume that we are working over a polynomial ring $\mathbb{K}[x_1, x_2, x_3, \ldots, x_n]$. Given as input a set of polynomials $\mathcal{F}$, for any $\epsilon > 0$, there is no polynomial-time algorithm $\mathcal{A}$ that*
- *Selects $\mathcal{F}' \subseteq \mathcal{F}$ with $|\mathcal{F}'| \geq (7/8 + \epsilon)|\mathcal{F}|$, and*
- *Returns a Gröbner Basis for $\mathcal{F}'$ with respect to any lexicographic order,*

unless $P = NP$.

Further, Theorem 7 holds as long as $\mathcal{F}$ is allowed to select polynomials from the core subset we identified in Theorem 6: restricting to polynomials of degree 3, or to polynomials that contain at most 3 variables each, etc will not ensure polynomial-time solvability.

The proof of Theorem 7 is similar to that for Theorem 2, omitting the second stage of truth assignment based on coin flipping. We construct the same polynomial system based on the clauses of a satisfiable instance of Max Non-Mixed Exactly 3SAT, apply the $\mathcal{A}$ described, perform elimination, and construct a partial truth assignment based on the solution a. This partial truth assignment already satisfies a $(7/8 + \epsilon)$-fraction of the clauses (since $\mathcal{A}$ returned a Gröbner basis corresponding to some subset of $(7/8 + \epsilon)$ of the clause polynomials), so any unassigned variables may be assigned truth values arbitrarily.

3. Robust hardness of graph-coloring

Our q-Fractional Gröbner Basis Problem considers a flavor of "approximate solution" that seems rather different from the most-widely-studied notion of an "approximate solution" in combinatorial optimization. Unlike a *β-approximation algorithm* that is tasked with delivering a solution of provably-good quality that is feasible for the full input, we insist instead on a solution of perfect quality [5] that need only be feasible for a fraction of the original input (where the algorithm may select *which fraction* of the input, subject to a few constraints). We feel this latter notion is quite natural for the Gröbner Basis Problem, but it is also a meaningful notion for traditional problems in combinatorial optimization. In this section, we consider a comparable problem in graph-coloring to the q-Fractional Gröbner Basis Problem. We point out that either positive algorithmic results or negative hardness results for this problem would immediately yield improved results for traditional questions about graph coloring.

Definition 8 (**q-Fractional Graph-Coloring Problem**)**.** Given as input a 3-colorable graph G on vertex set V, for specified $q \in [0, 1]$, output the following:
- $V_{in} \subseteq V$, such that $|V_{in}| \geq q|V|$.
- A 3-coloring of the induced subgraph for G on V_{in}.

It would be possible to define a similar problem in which a fraction of the edge set is taken, instead of the vertex set. However, such a problem is trivial even for $q = 1/2$. Namely, by greedily dividing the vertices of G into two parts, we can find a bipartite subgraph of G containing at least half the edges, which can be colored using two colors.

We now show that proving the q-Fractional Graph-Coloring Problem either to be in P or to be NP-hard would provide major contributions to the area of approximate

[5] That is, an actual Gröbner Basis.

graph-coloring. It was proven by Khanna et al. [17] and by Guruswami and Khanna [10] that it is NP-hard to color a 3-colorable graph in 4 colors. In fact, [10] shows something slightly stronger than this: even 4-coloring a 3-colorable graph so that nearly all edges are satisfied is NP-hard. However, it is not known whether one can efficiently color a 3-colorable graph in k colors, where k is any constant larger than 4. Under assumptions related to the Unique Games Conjecture, Dinur et al. [8] showed that it would be NP-hard to color a 3-colorable graph in any constant number of colors. Under similar assumptions, the statement was strengthened by Dinur and Shinkar, showing that it would be NP-hard to color a 4-colorable graph in $o(\log^c(n))$ colors for some constant $c > 0$.

A simple polynomial-time algorithm exists to color a 3-colorable graph in $O(n^{0.5})$ colors [25], where n is the number of vertices. Using semidefinite programming, this bound has been improved by [1, 2, 4, 14, 24]. Most recently, Kawarabayashi and Thorup made several advances by merging combinatorial and semidefinite approaches to achieve $O(n^{0.19996})$ (see [15, 16]).

Suppose first that there exists some polynomial time algorithm $\mathcal{A}$ for solving the q-Fractional Graph-Coloring Problem, for some value $q > 0$. Then, given a 3-colorable graph G on n vertices, we could color the entire graph in the following manner: Apply $\mathcal{A}$ to color at least qn vertices, so that at most $(1 - q)$ vertices are left uncolored. Now apply $\mathcal{A}$ once more to the remaining vertices using three new colors, so that at most $(1 - q)^2 n$ vertices are left uncolored. Continuing in this manner, we can apply $\mathcal{A}$ a total of $\log n / \log(1/(1-q))$ times to color all the vertices. This uses $O(\log n)$ colors and can be achieved in polynomial time, significantly improving upon existing algorithmic results.

Conversely, suppose that it is NP-hard to solve the q-Fractional Graph-Coloring Problem, for every value $q > 0$. Then, we argue that it must be NP-hard to k-color a 3-colorable graph, for every constant k. Suppose towards contradiction that some polynomial time algorithm existed. Taking $q = 3/k$, we could apply this algorithm and select the three most abundant colors within the coloring to form our vertex set V_{in}, leading to a contradiction.

4. The Strong c-Partial Gröbner Problem

Finally, we give a new result for a different previously-studied model of robust hardness for Gröbner basis computation. Given a system of polynomials $\mathcal{F}$ on a set of variables X, call $Y \subseteq X$ an *independent set of variables* provided that there exists no pair of variables in Y that appear in a common polynomial from $\mathcal{F}$. Recall the Strong c-Partial Gröbner Problem as defined by De Loera, et al. [7].

Definition 9. Given a set $\mathcal{F}$ of polynomials on a set X of variables, we say that $X' \subset X$ is an *independent set* of variables if these variables do not co-occur in any element of $\mathcal{F}$.

Definition 10 (from De Loera et al. [7])**.** Define the *Strong c-Partial Gröbner Problem* as follows. Given as input, a set $\mathcal{F}$ of polynomials on a set X of variables, output the following:

- disjoint $X_1, \ldots, X_b \subseteq X$, such that $b \leq c$ and each X_i is an independent set of variables,
- a Gröbner basis for $\langle \mathcal{F}_{in} \rangle$ over $X_{in} = X \backslash (\cup_{i=1}^{b} X_i)$ (i.e., we have taken away at most c independent sets of variables), where the monomial order of X is restricted to a monomial order on X_{in}.

De Loera, et al. [7] showed that for any system of polynomials subject to degree bound $k \geq 3$, unless $P = NP$, there is no polynomial-time algorithm for the Strong $(2\lfloor \frac{k}{3} \rfloor - 1)$-Partial Gröbner Problem. That is, even giving an algorithm the freedom to choose to ignore $(2\lfloor \frac{k}{3} \rfloor - 1)$ independent sets of variables (and all of the polynomials in which they appear) doesn't necessarily yield a polynomial system for which Gröbner basis computation is possible in polynomial time. Our Theorem 2 for the q-Fractional Gröbner Basis Problem gives a second way to understand robust hardness of Gröbner basis computation (the sets of generators which may be ignored under the q-fractional model have a different description). In both results, a surprising finding is that the robust hardness holds for polynomial systems with very low degree bounds (namely, degree of at most 3).

Now suppose that the number of independent sets of variables whose polynomials can be ignored by the algorithm is bounded by a much larger linear function of the degree bound k (rather than roughly $\frac{2}{3}k$ from [7]). Is computing a Gröbner basis for the remaining polynomial system still NP-hard? We prove the answer is always yes, provided that the degree bound allowed for the polynomial system is sufficiently large.

Precisely, we prove the following:

Theorem 11. *For every constant $h > 0$, there is a constant k_h such that the following problem is NP-hard. Given a polynomial system of maximum degree k_h, solve the Strong $(hk_h - 1)$-Partial Gröbner Problem for some lexicographic order.*

We prove Theorem 11 by reduction from a result on graph coloring due to Lund and Yannakakis [19].

Theorem 12 (Theorem 2.8 in [19])**.** *For every constant $g > 1$, there is a constant k_g such that the following problem is NP-hard. Given a graph G, distinguish between the case that G is colorable with k_g colors and the case that the chromatic number of G is at least gk_g.*

Proof of Theorem 11. Our proof is by contradiction. For arbitrary fixed $h > 0$, assume that for every possible k there exists a polynomial-time algorithm $\mathcal{A}_k$ to solve the Strong $(hk - 1)$-Partial Gröbner Problem.

Let $g = h + 1$. Then $g > 1$. Apply Theorem 12: there exists some constant k_g such that it is NP-hard to distinguish whether a graph is k_g-colorable or has chromatic number at least gk_g. Given a graph $G = (V, E)$ that is guaranteed to be one of these two types, we present a polynomial-time method to distinguish between the two. In particular, if G is k_g-colorable, then our method will produce a proper coloring of G by $gk_g - 1$ colors. If our method fails to produce such a coloring, then we will be forced to conclude that G is of the second type (that is, G has chromatic number at least gk_g). A description of this method follows.

Method to distinguish "G is k_g-colorable" from $\chi(G) \geq gk_g$:
Write the k_g-coloring ideal for G, $\mathcal{I}_{k_g}(G)$. This coloring ideal is a polynomial system with $|V|$ variables and $|V| + |E|$ polynomials of maximum degree k_g. From our initial assumption, there exists a polynomial-time algorithm $\mathcal{A}_{k_g}$ to solve the Strong $(hk_g - 1)$-Partial Gröbner Problem in polynomial systems of maximum degree k_g, so apply $\mathcal{A}_{k_g}$

to $\mathcal{I}_{k_g}(G)$. As described in De Loera et al. [7], if G is k_g-colorable then a solution to the Strong $(hk_g - 1)$-Partial Gröbner Problem can be converted in polynomial time to a partial proper coloring of the nodes of G by the k_gth roots of unity (using at most k_g colors). The still-uncolored nodes of G correspond to $hk_g - 1$ independent sets of nodes also identified by $\mathcal{A}_{k_g}$. Color each of these independent sets by one additional color. Thus, provided that G is k_g-colorable, this gives a total proper coloring of G by at most

$$k_g + hk_g - 1 = k_g + (g-1)k_g - 1 = gk_g - 1 \text{ colors.}$$

Such a proper coloring must be returned if G is k_g-colorable, and cannot possibly be returned if G has chromatic number at least gk_g. In this case we return "G is k_g-colorable."

Alternatively, suppose that we attempt to implement the above when G is not k_g-colorable. As G is guaranteed to be one of the two types, if G is not k_g-colorable, then G must have chromatic number at least gk_g. Thus, the procedure we have described above must necessarily fail to produce a $(gk_g - 1)$-coloring. Consider the implication for $\mathcal{I}_{k_g}(G)$: every set of $(hk_g - 1)$ independent sets of variables (and all polynomials in which any of these appear) which could possibly be ignored by $\mathcal{A}_{k_g}$ will give a remaining polynomial system whose Gröbner basis is generated by 1 (as the existence of any root would give rise to a total proper coloring of impossible size). Thus, if G is not k_g-colorable, then $\mathcal{A}_{k_g}$ returns $\langle 1 \rangle$ for the Strong $(hk_g - 1)$-Partial Gröbner Problem. If $\langle 1 \rangle$ is returned, we conclude: "$\chi(G) \geq gk_g$."

For arbitrary G, to distinguish between the two possible cases our method constructs a polynomially-sized input to $\mathcal{A}_{k_g}$, which from our initial assumption has polynomial running time. Therefore, in polynomial-time our method solves a problem we know to be NP-hard from Theorem 12. Thus (unless P=NP) we have a contradiction, and must reject our initial assumption. □

5. Future directions

In designing and improving algorithms for Gröbner basis computation, it is natural to hope that incorporating the generators in some clever order, or developing the polynomial consequences of the full set of generators gradually (perhaps by computing a reduced Gröbner basis iteratively as generators are incorporated) can lead to efficient methods for restricted classes of polynomial systems. Indeed, in computational tests, techniques of this flavor sometimes provide substantial speed-ups for sample instances. At a theoretical level, are useful practical techniques like these strong enough to overcome super-polynomial time running times? We hope that Theorem 2 provides some insight on methods that aim to circumvent the general super-polynomiality of Gröbner basis computation by addressing some large "simple part" of the polynomial system first. For such methods to work efficiently, it seems that powerful restrictions on the polynomial system will be required.

Empirical observation has suggested that Gröbner basis computation is sometimes particularly slow for lexicographic orders. Our theoretical results focus on these orders and deliver the message that even for highly-restricted polynomial systems, even approximate Gröbner basis computation for lexicographic orders isn't possible efficiently. Can similar robust hardness results be proved for a wider class of orders? If not, is there

a connection between theoretical results about robust hardness and orders that appear fastest in practice for popular algorithms?

Interpreting results on the robust hardness of Gröbner basis computation in terms of a graphical representation of polynomial system structure may lead to interesting directions in the study of what polynomial systems admit efficient computation of a corresponding Gröbner basis. For example, in one standard graphical representation (see [5] and [7]), variables are nodes and edges describe co-appearance within some polynomial: ignoring variables and the polynomials containing them (as considered in this paper) removes nodes and clique subgraphs which contain those nodes. Perhaps results about hardness persisting after such removals can yield insight towards characterizations of graphical structure that allow for efficient computation of a Gröbner basis. For example, the authors would be very interested to see more inquiry around the exciting results of Cifuentes and Parrilo [5].

Acknowledgements

The authors would like to thank David Cox and Jesús De Loera for constructive input on these ideas.

References

[1] A. Blum. New approximation algorithms for graph coloring. *J. ACM*, 41:470–516, 1994.

[2] A. Blum and D. Karger. An $\tilde{O}(n^{3/14})$ algorithm for 3-colorable graphs. *Information Processing Letters*, 61, 1997.

[3] B. Buchberger. A theoretical basis for the reduction of polynomials to canonical forms. *ACM SIGSAM Bulletin*, 10, 1976.

[4] E. Chlamtac. *Non-Local Analysis of SDP-Based Approximation Algorithms*. PhD thesis, Princeton University, 2009.

[5] D. Cifuentes and P. Parrilo. Exploiting chordal structure in polynomial ideals: a Gröbner bases approach. *arXiv:1411.1745*.

[6] D. Cox, J. Little, and D. O'Shea. *Ideals, Varieties and Algorithms*. Springer Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1992.

[7] J.A. De Loera, S. Margulies, M. Pernpeintner, E. Riedl, D. Rolnick, G. Spencer, D. Stasi, and J. Swenson. Graph-coloring ideals: Nullstellensatz certificates, Groebner bases for chordal graphs, and hardness of Groebner bases. *Proc. Int. Symp. on Symbolic and Algebraic Computation*, 2015.

[8] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. *SIAM J. Computing*, 39:843–873, 2009.

[9] M. Giusti. Some effectivity problems in polynomial ideal theory. *EUROSAM*, 1984.

[10] V. Guruswami and S. Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discrete Math.*, 18:30–40, 2004.

[11] V. Guruswami and S. Khot. Hardness of Max 3SAT with no mixed clauses. In *20th Annual IEEE Conference on Computational Complexity*, pages 154–162, 2005.

[12] J. Håstad. Some optimal inapproximability results. *Journal of theACM*, 48(4):798–859, 2001.

[13] D.S. Johnson. Approximation algorithms for combinatorial problems. *J. Comput. Syst. Sci.*, 9:256–278, 1974.

[14] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45:246–265, 1998.

[15] Ken-ichi Kawarabayashi and Mikkel Thorup. Combinatorial coloring of 3-colorable graphs. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 68–75, 2012.

[16] Ken-ichi Kawarabayashi and Mikkel Thorup. Coloring 3-colorable graphs with o(nˆ{1/5}) colors. In *31st International Symposium on Theoretical Aspects of Computer Science*, pages 458–469, 2014.

[17] S. Khanna, N. Linial, and S. Safra. On the hardness of approximating the chromatic number. *Proc. Israel Symp. on Theory and Computing Systems*, pages 250–260, 1993.

[18] K. Kühnle and E. W. Mayr. Exponential space computations of Gröbner bases. *Proc. Int. Symp. on Symbolic and Algebraic Computation*, 1996.

[19] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *J. ACM*, 41:960–981, 1994.

[20] E. W. Mayr. Some complexity results for polynomial ideals. *J. of Complexity*, 13:303–325, 1997.

[21] E.W. Mayr and S. Ritscher. Degree bounds for Gröbner bases of low-dimensional polynomial ideals. *Proc. Int. Symp. on Symbolic and Algebraic Computation*, 2010.

[22] M.H. Möller and T. Mora. Upper and lower bounds for the degree of Groebner bases. *EUROSAM*, 1984.

[23] S. Ritscher. Degree bounds for zero-dimensional Gröbner bases. Master's thesis, North Carolina State Univ., 2009.

[24] S. and E. Chlamtac. New approximation guarantee for chromatic number. *Proc. ACM Symp. on Theory of Computing*, 2006.

[25] A. Wigderson. Improving the performance guarantee for approximate graph coloring. *J. ACM*, 30:729–735, 1983.